16.6 Embedded SoC Resource Manager to Control Temperature and Data Bandwidth

Makoto Saen[1], Kenichi Osada[1], Satoshi Misaka[1], Tetsuya Yamada[1], Yoshitaka Tsujimoto[1], Yuki Kondoh[1], Tatsuya Kamei[2], Yutaka Yoshida[2], Ei Nagahama[2], Yusuke Nitta[2], Takayasu Ito[2], Tadashi Kameyama[2], Naohiko Irie[1]

[1]Hitachi, Tokyo, Japan
[2]Renesas Technology, Tokyo, Japan

Some systems have recently been integrated on a chip [1]; however, available resources, such as power consumption and data bandwidth, are severely limited and these limitatation will continue in the future [2]. The key to solve this problem is to develop a function to efficiently distribute limited resources to dozens of intellectual property blocks (IPs). Therefore, we developed an SoC resource manager responsible for supervising the whole SoC and allocating the limited resources among IPs. The resource manager controls the operation temperature and data bandwidth using a variety of monitored information: operating temperature, operating frequency of IPs, and the number of operations executed by IPs. The thermal control improves the reliability for real-time operation under various environment temperatures using an adaptive thermal control algorithm. The data bandwidth control allocates the bandwidth efficiently using a dynamic arbitration policy.

Figure 16.6.1 shows a block diagram of the chip. The resource manager controls the operation frequencies of the two processors (SH-CPU0/1) and the two hardware IPs (HWIP0/1) based on the operating temperature of the chip. Furthermore, the manager controls the data bandwidth based on the operating speed of IPs. The resource management modules consist of control parts and monitors. The control parts execute the thermal and data bandwidth control using the information obtained from the monitors. The control parts consist of a small processor core (RM core), a peripheral module (RM peripheral module), and a bus arbiter (RM arbiter) embedded in an on-chip bus. The monitors include a thermal sensor, an operation frequency table in a clock controller (CLKC), and a performance monitor (PMN) that gives information about the operation speed of the IPs to the RM core.

Thermal problems have already been identified in high performance processors [3]. Conventional thermal control reduces the operation frequency of a chip to low-power mode just before the operating temperature exceeds the device limit. However, embedded SoCs must execute real-time operations under various environment temperatures. This makes thermal control more difficult. The proper frequency for controlling the operating temperature changes depending on the temperature of the environment. Excessive slowdown of the operation frequency generates periods when real-time operations do not meet the performance requirements of the system. Therefore, the resource manager minimizes the slowdown of the operating frequency using a thermal control algorithm that adapts to the change in the environmental temperature.

The adaptive thermal control algorithm is based on the thermal model in Fig. 16.6.2. The SoC is modeled as a heat source with thermal capacitance, C_{si}, and thermal resistance, θ_{si}. The heat source consists of a heat source, H_{sw}, from the switching power, P_{sw}, and a heat source, H_{lk}, from the leakage power, P_{lk}. The thermal resistance of the package is θ_{pkg}. The environmental temperature is T_a. A thermal sensor measures the operating temperature, T_j, of the SoC. Since T_a varies and cannot be directly observed, it prevents proper control of the frequency. In addition, since the leakage power P_{lk} is dependent on the operating temperature, P_{lk} needs to be monitored. Therefore, in the first stage of the algorithm (Fig. 16.6.2), T_a and H_{lk} are calculated based on the P_{sw} and the change in T_j. Results from the first stage are used in the second stage to calculate the limit of switching power, P_{swmax}, required to control T_j to a target temperature, T_{jT}. The RM core allocates the calculated P_{swmax} to IPs based on a switching power distribution table that defines the distribution to IPs for each P_{swmax}.

A block diagram of the thermal control is shown in Fig. 16.6.3. There is the RM core for executing the algorithm, a thermal sensor for finding T_j, and an operating frequency table in a clock controller (CLKC) for finding P_{sw}. In addition, the RM peripheral module sends the thermal measurement by the thermal sensor to the RM core. To reduce the power consumption overhead brought by the thermal control, the RM peripheral module has two features. One is to use interrupt signals to send the thermal information to the RM core. The other is a filter block to remove excess interrupts cased by spike noise and fluctuations detected in the temperature levels of the interrupt.

The resource manager also controls the data bandwidth of each IP using the RM arbiter, a timer, and a performance monitor (Fig. 16.6.4). This bandwidth control allocates the bandwidth to IPs, depending on the actual needed bandwidth, which changes dynamically. The RM core calculates the priority of data-access requests by each IP based on the difference between the actual operation speed and the target operation speed. The operation speed is defined as the number of operations executed in a given time period. The slower the actual operation speed than the target speed, the higher the priority.

The resource manager is fabricated in 90nm CMOS (Fig. 16.6.5). The frequency of the RM core and peripheral module is 50MHz. The total area of the RM core and peripheral module is 0.4 mm^2, and the area of the thermal sensor is 0.3mm^2. Figure 16.6.6 shows measurement results of the thermal sensor block.

Figure 16.6.7 shows the simulation results used to evaluate the thermal control. We also develop a simulator to calculate the relationship between the power consumption, temperature, and IP operation. The upper graph plots the dependence of T_a on the upper limit frequency to control the operating temperature under 125°C. For example, using the resource manager, at a T_a of 75°C, SH-CPU operation frequencies of 600 and 300MHz are obtained. In this case, the total operation frequency is 900MHz. Here, we define conventional control as the simple control in which the operation clock is stopped if the operation temperature exceeds a temperature limit. So, the control generates a period when no operations are executed at a given temperature. If a target system requires total operation frequency of 600MHz for real-time operations, the resource manager allows a T_a of about 85°C. The results indicate that the manager improves the upper limit of T_a by about 30°C.

Acknowledgements:
We thank H. Shimizu, O. Mouri, T. Hattori, N. Mitsuishi, M. Hiraki, A. Hasegawa, and all members of the development team for fruitful discussion and technical support.

References:
[1] T. Hattori, T. Irita, M. Ito, et al., “A Power Management Scheme Controlling 20 Power Domains for a Single-Chip Mobile Processor,” *ISSCC Dig. Tech. Papers*, pp. 542-543, Feb., 2006.
[2] S. Torii, S. Suzuki, H. Tomonaga, et al., “A 600MIPS 120mW 70mA Leakage Triple-CPU Mobile Application Processor Chip,” *ISSCC Dig. Tech. Papers*, pp. 136-137, Feb., 2005.
[3] C. Poirier, R. McGowen, R. Bostak, and S. Naffziger, “Power and Temperature Control on a 90nm Itanium-Family Processor,” *ISSCC Dig. Tech. Papers*, pp. 304-305, Feb., 2005.

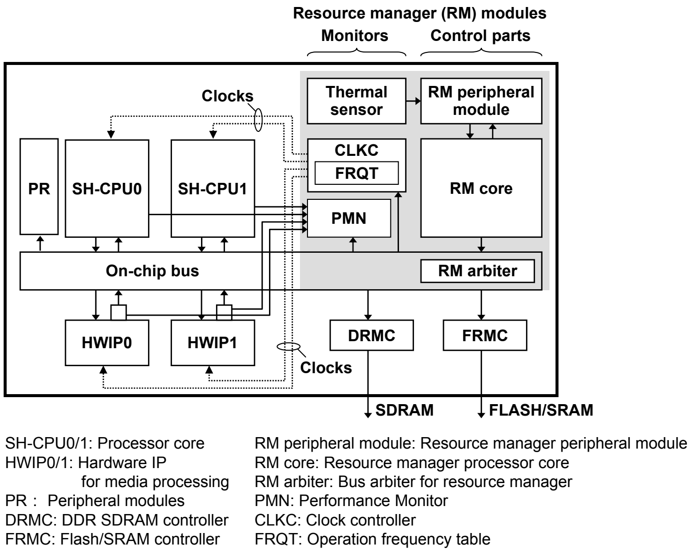

Figure 16.6.1: Block diagram of chip.

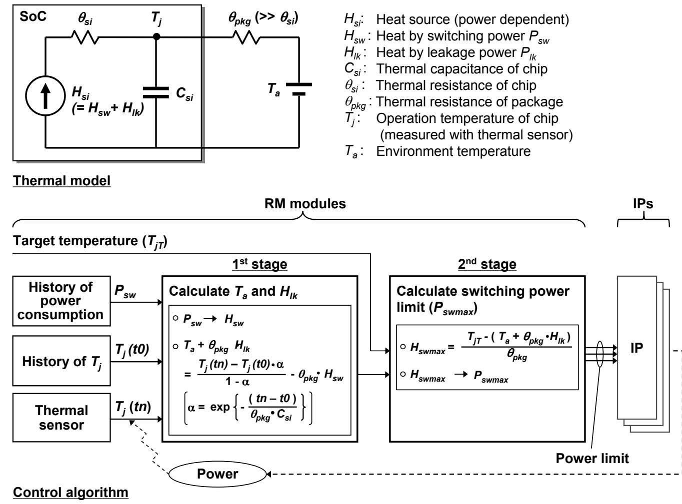

Figure 16.6.2: Thermal control algorithm.

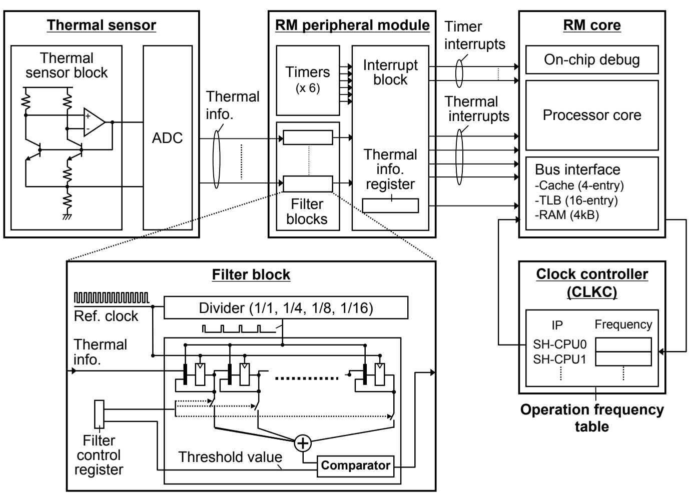

Figure 16.6.3: Architecture for thermal control.

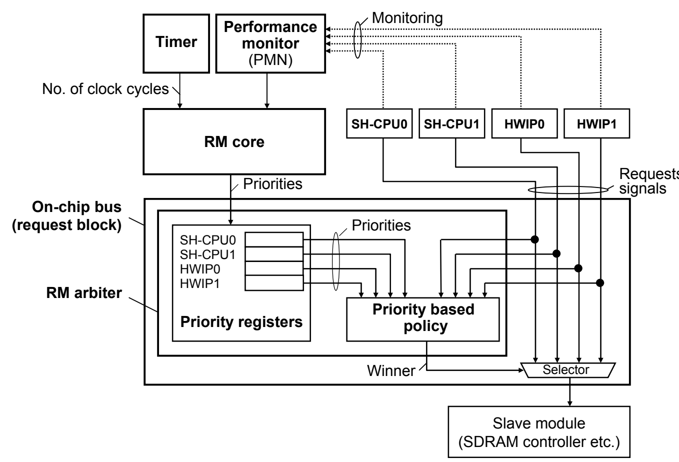

Figure 16.6.4: Block diagram of data bandwidth control.

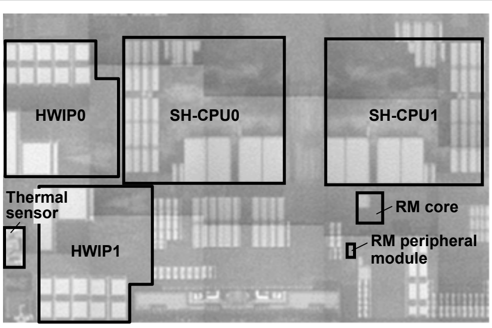

Technology process		90 nm, 8-metal
Frequency	RM core	50 MHz
	RM peripheral module	50 MHz
	On-chip bus	300 MHz
	SH-CPU0, 1	600 MHz
Area	RM core RM peripheral module	0.4 mm²
	Thermal sensor	0.3 mm²

Figure 16.6.5: Experimental chip micrograph.

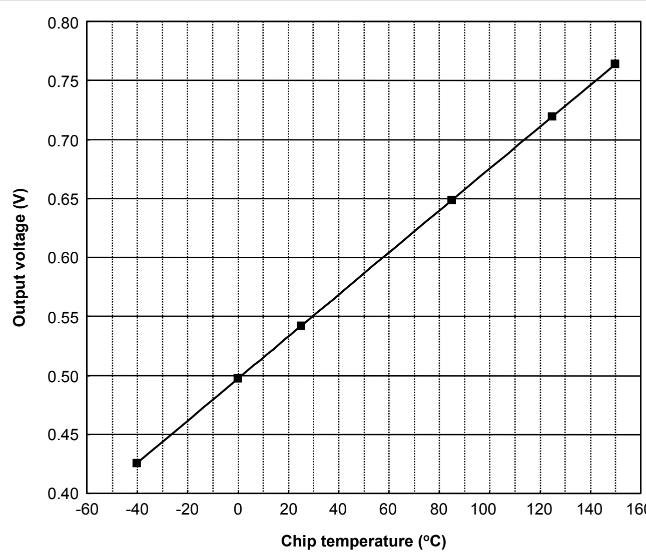

Figure 16.6.6: Measurement results of thermal sensor block.

Continued on Page 604

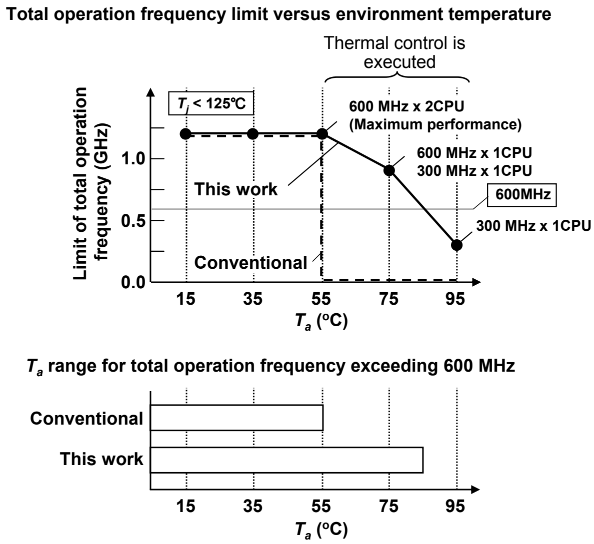

Figure 16.6.7: Effect of adaptive thermal control (simulation results).